Logical foundations of categorization theory
Lecture 1

Fei Liang and Alessandra Palmigiano

July 26, 2021

Abstract

In the first part this lecture, we discuss the extant foundational views on categorization theory, and their pros and cons. In the second part, we give an informal outline of the extant mathematical models for categorization and to which extent they reflect insights from the foundational views. In the third part, we argue that there is a mathematical approach that unifies these various models, and – time permitting –we start giving some details about this approach.

1 Introduction

Categories. Categories are cognitive tools that humans use to organize their experience, understand and function in the world, and understand and interact with each other, by grouping things together which can be meaningfully compared and evaluated. Categorization is the basic operation humans perform e.g. when they relate experiences/actions/objects in the present to ones in the past, thereby recognizing them as instances of the same type. This is what we do when we try and understand what an object is or does, or what a situation means, and when we make judgments or decisions based on experience. Categorization is the single cognitive mechanism underlying *meaning*-attribution, *value*-attribution and *decision-making*. In this course, we will introduce the basic facts and properties (cf. [CFPPTW16, CFPPTW17]) of logical systems describing and reasoning about categories and categorization, thereby laying the ground of a mathematical-logical environment in which these three cognitive processes can be formally analyzed in their relationships to one another.

Relevance. Nowadays, categories are key to the theories and methodologies of a wide range of fields in the social sciences and AI: in cognitive anthropology, cultures are analysed and compared in terms of their *categorization behaviour* [d'A95]; in linguistics, categories are central to the mechanisms of *grammar generation* [Cr91]; in cognitive science, *categorical perception* [Ha87, Be03] has been identified as the most basic cognitive activity, and used to analyze higher-order cognitive activities; in AI, classification techniques are core to *pattern recognition*, *data mining*, *text mining*, and *knowledge discovery* in databases; in psychology, the *dynamics of category-acquisition* serve to illuminate the relationship between language and thought [VHV12]; in sociology, categories are used to explain the construction of *social identity* [Je00, Je14] and the organization of experience (cf. *frame analysis* [Go74]); in management science, categories are used to predict how products and producers will be *perceived* and *evaluated* by consumers and investors [Wi11, CPT17, PD16, WLJ14, DK17]. Therefore, developing the basis for a logical theory of categorization will also create a *lingua franca* which will make it possible to establish novel connections and transfer results and insights between these fields.

2 Extant foundational approaches

The literature on the foundations of categorization theory (cf. [CL05]) displays a variety of definitions, theories, models, and methods, each with its pros and cons. Below we briefly review the most influential approaches.

The classical theory. The classical theory of categorization [SM81] goes back to Aristotle, and is based on the insight that all members of a category share some fundamental features which define their membership. Accordingly, categorization is viewed as a deductive process of reasoning with necessary and sufficient conditions, resulting in categories with sharp boundaries, which are represented equally well by any of their members. The classical view has inspired influential approaches in machine learning such as conceptual clustering [Mi80,Fi87]. However, this view runs into difficulties when trying to accommodate a *new* object or entity which would intuitively be part of a given category but does not share all the defining features of the category. Other difficulties are:
(a) providing an exhaustive list of defining features (Wittgenstein: what is a game?);
(b) how to deal with unclear cases (is it blue or is it green?);
(c) the existence of members of given categories which are judged to be better representatives of the whole class than others (if I ask you to think of a bird, would you think of a sparrow or would you think of a penguin?).
All these issues motivated the introduction of the second view.

The prototype view. This view was developed by Rosch and Lakoff [La99,Ro05]. According to it, categorization is not a deductive process, but is rather an *inductive* process by which we try and find the best match between the features of an object and those of the closest prototype(s). In this way, membership in a category does not need anymore to be decided through the satisfaction of an exhaustive list of features; this approach allows for unclear cases (the closer an object is to the prototype, the easier it is to decide whether it is a member, and borderline cases are also possible), and embracing the empirically verified intuition that people regard membership in most categories as a matter of degrees, and certain members as more central (or prototypical) to a category than others. However, how do we generate prototypes in our minds? Prototype theory does not have an answer to this issue.

The exemplar view. To address this issue, the exemplar theory [SM02] was proposed, according to which individuals make category judgments by comparing new stimuli with instances already stored in memory (the "exemplars"). However, the existence of instances or prototypes of a given category presupposes that this category has already been defined, hence both the prototype and the exemplar view run into a circularity problem. Moreover, it has been argued that similarity-based theories of categorization such as exemplar theory and prototype theory fail to address the problem of explaining 'why we have the categories we have', or, in other words, why certain categories seem to be more cogent and coherent than others. Even more fundamentally, similarity might be imposed rather than discovered (do things belong in the same category because they are similar, or are they similar because they belong in the same category?), i.e. might be the effect of conceptual coherence rather than its cause.

The theory-based view. Pivoting on the notion of coherence for category-formation, the theory-based view on categorization [MM85] posits that categories arise in connection with theories (broadly understood so as to include also informal explanations). For instance, ice, water and steam can be grouped together in the same category on the basis of the theory of phases in physical chemistry. The coherence of categories proceeds from the coherence of the theories on which they are based. This view of categorization allows one to group together entities which would be scored as dissimilar using different methods; for instance, it allows to group together a gold watch, the school report of one's grandfather, and the naked ownership of a piece of land in the category of "things one wants one's children to inherit", which is based on one's theory of what family is. However, the theory-based view does not account for the intuition that categories themselves are the building blocks of theory-formation, which again results in a circularity problem, and does not account for how changes in the theories account for changes in the categories.

Summing up. Although each of these views provides useful insights on the essence of categories, none of them satisfactorily solves all the issues and provides an overarching approach to the foundations of categorization capable to reconcile the coexistence of seemingly contradictory aspects, such as the sharp vs vague nature of categorical boundaries, or accounting for categorical coherence and stability in the face of contextual changes.

Our suggested approach. Notwithstanding its problems, the only theory of categorization which deserves to be called 'theory' still remains the classical one. So, rather than try and replace the classical theory with another one, we should endow the classical theory with *improved formal tools* so as to be able to account for the aspects of categories that are not yet covered by the classical theory.

3 Extant mathematical approaches

In the extant literature, there are four most influential mathematical approaches to the representation of categories and concepts. Below, we review these approaches while simultaneously discussing to which extent they can be linked to the extant foundational approaches above, and to which extent they have been linked to logic.

Formal Concept Analysis (FCA). This is a method of data analysis pioneered by Wille [GW12] and based on Birkhoff's representation theory of complete lattices [Bi40]. In FCA, databases are represented as *formal contexts*, i.e. structures (A,X,I) such that A and X are sets, and $I \subseteq A \times X$ is a binary relation. Intuitively, A is understood as a collection of objects, X as a collection of features, and for any object a and feature x, the tuple (a,x) belongs to I exactly when object a has feature x. Every formal context (A,X,I) can be associated with the collection of its *formal concepts*, i.e. the tuples (B,Y) such that B is a subset of A, Y is a subset of X, and $B \times Y$ is a rectangle maximally included in I. The set B is the *extension* of the formal concept (B,Y), and Y is its *intension*. Because of maximality, the extension of a formal concept uniquely identifies and is identified by its intension. Formal concepts can be partially ordered; namely, (B,Y) is a *subconcept* of (C,Z) exactly when B is a subset of C, or equivalently, when Z is a subset of Y. Ordered in this way, the concepts of a formal context form a complete lattice (i.e., the least upper bound and the greatest lower bound of every set of formal concepts exist), and by Birkhoff's theorem, every complete lattice is isomorphic to a concept lattice. FCA provides a very elegant and successful mathematical representation of key insights into categories and concepts; for instance, while many approaches identify concepts with their extension, in FCA, intension and extension of a context are treated on a par, i.e., the intension of a concept is just as essential as its extension. While FCA has tried to connect itself with various cognitive and philosophical theories of concept formation, it is most akin to the classical view.

Conceptual spaces. The second mathematical approach to the representation of categories and concepts was introduced by Gärdenfors and employs conceptual spaces [Gä04]. These are multi-dimensional geometric structures, the components of which (the quality dimensions) are intended to represent basic features–e.g. color, pitch, temperature, weight, time, price–by which objects (represented as points in the product space of these dimensions) can be meaningfully compared. Each dimension is endowed with its appropriate geometric (e.g. metric, topological) structure. Concept-formation in conceptual spaces is modelled according to a similarity-based view of concepts. Specifically, if each dimension of a conceptual space has a metric, these metrics translate in a notion of distance between the objects represented in the space, which models their similarity, so that the closer their distance, the more similar they are. Concepts are represented as convex sets of the conceptual space (a subset is convex if it includes the segments between any two of its points. In the Euclidian plane, squares are convex while stars are not). The geometric center of any such concept is a natural interpretation of the prototype of that concept. Conversely, any finite set of points (understood as the set of prototypes) gives rise to a tiling of the space, in which each point is assigned to its closest prototype(s). If the conceptual space is endowed with a Euclidian metrics, the tiles so generated are convex sets. Although conceptual spaces have been argued to support both non-monotonic reasoning [Gä04] and fuzzy reasoning [Go05], neither view has been concretely realized.

Vector space models. The third approach, making use of vector space models (VSMs) [TP10, Cl15], is technically akin to Gärdenfors' approach, although it is motivated differently. This approach was developed in the field of information retrieval and then imported to natural language processing and computational linguistics. Given a corpus of n documents, words are modelled as the n-tuples (vectors) of the frequencies with which they occur in each document. While VSMs are used in information retrieval to extract information on what documents are about on the base of word-frequencies, in linguistics they are used to extract information about the similarity of meaning of words, based on their co-occurrences in a corpus of documents. The motivating insight for the use of VSMs in linguistics is the distributional hypothesis [Wi10], which posits that words occurring in similar contexts are likely to have similar meanings. The natural notion of distance between vectors allows to quantify the similarity of meaning of two given words (the closer their distance as vectors, the more similar their meaning). VSMs are widely considered the most successful approach to lexical semantics, and recently, this approach has been augmented so as to make it compatible with a compositional theory of meaning [CSC10, GS11, GSCP14], which is an important step in the direction of logic. The applications of VSMs focus mainly on similarity judgments rather than higher-level concepts, and consequently, VSMs have not been explicitly related with any theory of concepts, although, mathematically, they have been recently shown to be a specialization of conceptual spaces [MAPW15], and seminal steps have been taken [BCLM16] to endow this framework with a reasoning machinery accounting for concept-entailment.

Event spaces The fourth approach makes use of event spaces in conceptual clustering [Mi80]. These structures are cartesian products of a finite number of (typically finite, discrete or linearly ordered) domains of interpretation of a variable. Objects (events) are represented as the tuples of the values they score on each variable, and several natural distance notions can be introduced to capture different kinds of graded similarity between events. Unlike the other mathematical approaches discussed above, a logical perspective is explicitly introduced in terms of a quantifier-free predicate logical language which is defined for each event space. The atomic formulas of this language, called selectors, are (essentially) membership statements of the form $x \in X$ for some subset X of the domain of interpretation of the variable x. The original presentation of conceptual clustering is inspired by the classical view on categorization, and this is witnessed by the fact that concepts are represented syntactically as conjunctions of selectors, i.e. in terms of all the properties that members of the category need to satisfy. However, subsequent variations (e.g. [Mi83]) draw insights from the theory-based view. From a technical viewpoint, associating a logical language with each event space only serves to provide an a priori specification of a finite list of subsets of the event space embodying the relevant concepts, which could be otherwise specified in purely set-theoretic, non-logical terms, and indeed, this set-theoretic perspective on concepts is the dominant one pursued in conceptual clustering.

4 Unification via closure operators

The unification that we propose starts with the technical insight that the four mathematical approaches described above can be unified by a single notion – that of *closure operator*, cf. [DP95] – which accounts for the concept-formation in each approach. If $(P,\leq)$ is a partially ordered set, a map $c : P \to P$ is a *closure operator* on $(P,\leq)$ if, for all x,y in P:
(a) $x \leq c(x)$;
(b) if $x \leq y$ then $c(x) \leq c(y)$;
(c) $c(c(x)) = c(x)$.
Given a topological space X, the map associating any subset Y with the smallest closed set containing Y is a closure operator on the powerset of the domain of X. Given a logic L, the map associating any set of formulas Γ with the set of formulas which are logical consequences of Γ is a closure operator on the powerset of the formulas of L. In a Euclidean space X, the map associating any subset Y with its convex hull (i.e. the smallest convex set containing Y) is a closure operator on the powerset of X. Any closure operator on $(P,\leq)$ is completely determined by the set $\{x \in P \mid x = c(x)\}$ of its closed elements. The set of closed subsets of any closure operator on the powerset of a set is a complete lattice, with greatest lower bounds given by intersections, and least upper bounds given by closures of unions. Conversely, every complete lattice is isomorphic to the lattice of closed sets of some closure operator on the powerset of some set (this is an equivalent restatement of Birkhoff's representation theorem). The concept lattice associated with a formal context (A,X,I) is isomorphic to the lattice of closed sets of the closure operator $c : P(A) \to P(A)$ defined by mapping any subset B of A to its unique superset C such that C is the "side" (the first projection) of some maximal rectangle $C \times Z$ included in I. Equivalently, the same concept lattice is (dually) isomorphic to the lattice of closed sets of the closure operator $c : P(X) \to P(X)$ defined by mapping any subset Y of X to its unique superset Z such that Z is the "side" (the second projection) of some maximal rectangle $C \times Z$ included in I. The lattice of concepts of a conceptual space (or vector space) X is the lattice of closed sets of the closure operator $c : P(X) \to P(X)$ defined by mapping any subset Y of X to its convex hull. Finally, the semantic concepts of event spaces are generated by taking all intersections of a set of basic concepts. This gives rise to a collection of subsets endowed by construction with the characterizing property (closure under arbitrary intersection) of the set of closed sets of a closure operator.

Upshot. The discussion above can be summarized as follows: the four mathematical approaches to categorization model concept-generation by means of a single underlying mechanism, captured by the (order-theoretic) notion of closure operator. Closure operators underlie both the generation of logical theories from sets of formulas, and the generation of categories and concepts from sets of elements. Closure operators and complete lattices are intimately related via Birkhoff's representation theorem. This relation gives a mathematical backbone to the intuition that categories do not arise in isolation, but as elements of hierarchies of categories (modelled as complete lattices).

References

[BCLM16] Bankova, D., Coecke, B., Lewis, M., & Marsden, D. (2016). Graded entailment for compositional distributional semantics. arXiv preprint arXiv:1601.04908.

[Be03] Beer, Randall D. "The dynamics of active categorical perception in an evolved model agent." Adaptive Behavior 11.4 (2003): 209-243.

[Bi40] Birkhoff, Garrett. Lattice theory. Vol. 25. American Mathematical Soc., 1940.

[CPT17] Cattani, Gino, Joseph F. Porac, and Howard Thomas. "Categories and competition." Strategic Management Journal 38.1 (2017): 64-92.

[Cl15] Clark, Stephen. "Vector space models of lexical meaning. Handbook of Contemporary Semantics." (2015).

[CSC10] Coecke, Bob, Mehrnoosh Sadrzadeh, and Stephen Clark. "Mathematical foundations for a compositional distributional model of meaning." arXiv preprint arXiv:1003.4394 (2010).

[CL05] Cohen, Henri, and Claire Lefebvre, eds. Handbook of categorization in cognitive science. Elsevier, 2005.

[CFPPTW16] Conradie, W., S. Frittella, A. Palmigiano, M. Piazzai, A. Tzimoulis, N. Wijnberg. "Categories: How I Learned to Stop Worrying and Love Two Sorts." Proc. WoLLIC 2016, LNCS 9803, pp. 145-164, 2016. ArXiv preprint 1604.00777.

[CFPPTW17] Conradie, W., S. Frittella, A. Palmigiano, M. Piazzai, A. Tzimoulis, N. Wijnberg. "Toward an Epistemic-Logical Theory of Categorization". Proc. TARK 2017, J.Lang ed, EPTCS 251, pp. 167-186, 2017.

[Cr91] Croft, William. Syntactic categories and grammatical relations: The cognitive organization of information. University of Chicago Press, 1991.

[d'A95] d'Andrade, Roy G. The development of cognitive anthropology. Cambridge University Press, 1995.

[DP95] Davey, Brian A., and Hilary A. Priestley. Introduction to lattices and order. Cambridge university press, 2002.

[DK17] Durand, Rodolphe, and Mukti Khaire. "Where Do Market Categories Come From and How? Distinguishing Category Creation From Category Emergence." Journal of Management 43.1 (2017): 87-110.

[Fi87] Fisher, Douglas H. "Knowledge acquisition via incremental conceptual clustering."Machine learning 2.2 (1987): 139-172.

[Gä04] Gärdenfors, Peter. Conceptual spaces: The geometry of thought. MIT press, 2004.

[GW12] Ganter, Bernhard, and Rudolf Wille. Formal concept analysis: mathematical foundations. Springer Science & Business Media, 2012.

[Go74] Goffman, Erving. Frame analysis: An essay on the organization of experience. Harvard University Press, 1974.

[Go05] Goguen, Joseph. "What is a concept?." International Conference on Conceptual Structures. Springer Berlin Heidelberg, 2005.

[GS11] Grefenstette, Edward, and Mehrnoosh Sadrzadeh. ”Experimental support for a categorical compositional distributional model of meaning.” Proceedings of the Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics, 2011.

[GSCP14] Grefenstette, E., Sadrzadeh, M., Clark, S., Coecke, B., Pulman, S. (2014). Concrete sentence spaces for compositional distributional models of meaning. In Computing meaning (pp. 71-86). Springer Netherlands.

[Ha87] Harnad, S. (ed.) Categorical Perception: The Groundwork of Cognition. Cambridge Univ. Press, New York, 1987.

[Ha68] Harris, Zellig. ”Mathematical structures of language.” (1968).

[Je00] Jenkins, Richard. ”Categorization: Identity, social process and epistemology.” Current sociology 48.3 (2000): 7-25.

[Je14] Jenkins, Richard. Social identity. Routledge, 2014.

[KW10] Khaire, Mukti, and R. Daniel Wadhwani. ”Changing landscapes: The construction of meaning and value in a new market category-Modern Indian art.” Academy of Management Journal 53.6 (2010): 1281-1304.

[La99] Lakoff, George. ”Cognitive models and prototype theory.” Concepts: Core Readings (1999): 391-421.

[La90] Lancaster, Kelvin. ”The economics of product variety: A survey.” Marketing Science 9.3 (1990): 189-206.

[MAPW15] McGregor, S., Agres, K., Purver, M., and Wiggins, G. A. (2015). From distributional semantics to conceptual spaces: A novel computational method for concept creation. Journal of Artificial General Intelligence, 6(1), 55-86.

[Mi80] Michalski, Ryszard S. ”Knowledge acquisition through conceptual clustering: A theoretical framework and an algorithm for partitioning data into conjunctive concepts.” Journal of Policy Analysis and Information Systems 4.3 (1980): 219-244.

[Mi83] Michalski, Ryszard S. ”A theory and methodology of inductive learning.” Artificial intelligence 20.2 (1983): 111-161.

[MS83] Michalski, Ryszard S., and Robert E. Stepp. ”Learning from observation: Conceptual clustering.” Machine learning. Springer Berlin Heidelberg, 1983. 331-363.

[MM85] Murphy, Gregory L., and Douglas L. Medin. ”The role of theories in conceptual coherence.” Psychological review 92.3 (1985): 289.

[PD16] Paolella, Lionel, and Rodolphe Durand. ”Category spanning, evaluation, and performance: Revised theory and test on the corporate law market.” Academy of Management Journal 59.1 (2016): 330-351.

[PB15] Pontikes, Elizabeth G., and William P. Barnett. ”The Persistence of Lenient Market Categories.” Organization Science 26.5 (2015): 1415-1431.

[Ro05] Rosch, Eleanor. ”Principles of categorization.” Etnolingwistyka. Problemy jezyka i kultury 17 (2005): 11-35.

[SM81] Smith, Edward E., and Douglas L. Medin. Categories and concepts. Cambridge, MA: Harvard University Press, 1981.

[SM02] Smith, Edward E., and Douglas L. Medin. ”The exemplar view.” Foundations of cognitive psychology: Core readings (2002): 277-292.

[TP10] Turney, Peter D., and Patrick Pantel. ”From frequency to meaning: Vector space models of semantics.” J. of artificial intelligence research 37 (2010): 141-188.

[VHV12] Vygotski, L. S., E. Hanfmann, and G. Vakar. Thought and language. MIT press, 2012.

[Wi11] Wijnberg, Nachoem M. ”Classification systems and selection systems: The risks of radical innovation and category spanning.” Scandinavian Journal of Management 27.3 (2011): 297-306.

[Wi10] Wittgenstein, Ludwig. Philosophical investigations. John Wiley & Sons, 2010.

[WLJ14] Wry, Tyler, Michael Lounsbury, and P. Devereaux Jennings. ”Hybrid vigor: Securing venture capital by spanning categories in nanotechnology.” Academy of Management Journal 57.5 (2014): 1309-1333.